

BET: Bayesian Ensemble Trees for Clustering and Prediction in Heterogeneous Data

Leo L. Duan ¹, John P. Clancy ² and Rhonda D. Szczesniak ^{3 4}

Summary

We propose a novel “tree-averaging” model that utilizes the ensemble of classification and regression trees (CART). Each constituent tree is estimated with a subset of similar data. We treat this grouping of subsets as Bayesian ensemble trees (BET) and model them as an infinite mixture Dirichlet process. We show that BET adapts to data heterogeneity and accurately estimates each component. Compared with the bootstrap-aggregating approach, BET shows improved prediction performance with fewer trees. We develop an efficient estimating procedure with improved sampling strategies in both CART and mixture models. We demonstrate these advantages of BET with simulations, classification of breast cancer and regression of lung function measurement of cystic fibrosis patients.

KEY WORDS: Bayesian CART; Dirichlet Process; Ensemble Approach; Heterogeneity; Mixture of Trees.

¹Department of Mathematical Sciences, University of Cincinnati

²Division of Pulmonary Medicine, Cincinnati Children’s Hospital Medical Center

³Division of Biostatistics and Epidemiology, Cincinnati Children’s Hospital Medical Center

⁴Corresponding author. Address: 3333 Burnet Ave, MLC 5041, Cincinnati, OH 45229. Phone:(513)803-0563, email: rhonda.szczesniak@cchmc.org

1 Introduction

Classification and regression trees (CART) (Breiman, Friedman, Olshen, and Stone, 1984) is a nonparametric learning approach that provides fast partitioning of data through the binary split tree and an intuitive interpretation for the relation between the covariates and outcome. Aside from simple model assumptions, CART is not affected by potential collinearity or singularity of covariates. From a statistical perspective, CART models the data entries as conditionally independent given the partition, which not only retains the likelihood simplicity but also preserves the nested structure.

Since the introduction of CART, many approaches have been derived with better model parsimony and prediction. The Random Forests model (Breiman, 2001) generates bootstrap estimates of trees and utilizes the bootstrap-aggregating (“bagging”) estimator for prediction. Boosting (Friedman, 2001, 2002) creates a generalized additive model of trees and then uses the sum of trees for inference. Bayesian CART (Denison, Mallick, and Smith, 1998; Chipman, George, and McCulloch, 1998) assigns a prior distribution to the tree and uses Bayesian model averaging to achieve better estimates. Bayesian additive regression trees (BART, Chipman, George, McCulloch, et al. (2010)) combine the advantages of the prior distribution and sum-of-trees structure to gain further improvement in prediction.

Regardless of the differences in the aforementioned models, they share one principle: multiple trees create more diverse fitting than a single tree; therefore, the combined information accommodates more sources of variability from the data. Our design follows this principle.

We create a new ensemble approach called the Bayesian Ensemble Trees (BET) model, which utilizes the information available in the subsamples of data. Similar to Random Forests, we hope to use the average of the trees, of which each tree achieves an optimal fit without any restraints. Nonetheless, we determine the subsamples through clustering rather than bootstrapping. This setting automates the control of the number of trees and also adapts the trees to possible heterogeneity in the data.

In the following sections, we first introduce the model notation and its sampling algorithm. We illustrate the clustering performance through three different simulation settings. We demonstrate the new tree sampler using homogeneous example data from a breast cancer study. Next we benchmark BET against other tree-based methods using heterogeneous data on lung function collected on cystic fibrosis patients. Lastly, we discuss BET results and possible extensions.

2 Preliminary Notation

We denote the i th record of the outcome as Y_i , which can be either categorical or continuous. Each Y_i has a corresponding covariate vector \mathbf{X}_i .

In the standard CART model, we generate a binary decision tree \mathbf{T} that uses only the values of \mathbf{X}_i to assign the i th record to a certain region. In each region, elements of Y_i are identically and independently distributed with a set of parameters $\boldsymbol{\theta}$. Our goals are to find the optimal tree \mathbf{T} , estimate $\boldsymbol{\theta}$ and make inference about an unknown Y_s given values of \mathbf{X}_s , where s indexes the observation to predict.

We further assume that $\{Y_i, \mathbf{X}_i\}$ is from one of (countably infinitely) many trees $\{\mathbf{T}_j\}_j$. Its true origin is only known up to a probability w_j from the j th tree. Therefore, we need to estimate both \mathbf{T}_j and w_j for each j . Since it is impossible to estimate over all j 's, we only calculate those j 's with non-negligible w_j , as explained later.

3 Bayesian Ensemble Tree (BET) Model

We now formally define the proposed model. We use $[\cdot]$ to denote the probability density. Let $[Y_i|\mathbf{X}_i, \mathbf{T}_j]$ denote the probability of Y_i conditional on its origination is from the j th tree. The mixture likelihood can be expressed as:

$$[Y_i|\mathbf{X}_i, \{\mathbf{T}_j\}_j] = \sum_j^\infty w_j [Y_i|\mathbf{X}_i, \mathbf{T}_j] \quad (1)$$

where the mixture weight vector has an infinite-dimension Dirichlet distribution with precision parameter α : $\mathbf{W} = \{w_j\}_j \sim \text{Dir}_\infty(\alpha)$. The likelihood above corresponds to $Y_i \stackrel{iid}{\sim} DP(\alpha, G)$, where DP stands for the Dirichlet process and the base distribution G is simply $[Y_i|\mathbf{T}_j]$.

3.1 Hierarchical Prior for \mathbf{T}

We first define the nodes as the building units of a tree. We adopt the notation introduced by Wu, Tjelmeland, and West (2007) and use the following method to assign indices to the nodes: for any node k , two child nodes are indexed as left ($2k+1$) and right ($2k+2$); the root node has index 0. The parent of any node $k > 0$ is simply $\lfloor \frac{k-1}{2} \rfloor$, where $\lfloor \cdot \rfloor$ denotes the integer part of a non-negative value. The depth of a node i is $\lfloor \log_2(k+1) \rfloor$.

Each node can have either zero (not split) or two (split) child nodes. Conditional on the parent node being split, we use $s_k = 1$ to denote one node being split (or interior), $s_k = 0$ otherwise (or leaf). Therefore, the frame of a tree with at least one split is:

$$[s_0 = 1] \prod_{k \in I_k} [s_{2k+1} | s_k = 1] [s_{2k+2} | s_k = 1]$$

where $I_k = \{k : s_k = 1\}$ denotes the set of interior nodes.

Each node has splitting thresholds that correspond to the m covariates in \mathbf{X} . Let the m -dimensional vector \mathbf{t}_k denote these splitting thresholds. Also, it has a random draw variable c_k from $\{1, \dots, m\}$. We assume s_k, \mathbf{t}_k, c_k are independent.

For the c_k th element $X_{(c_k)}$, if $X_{(c_k)} < t_{(c_k)}$, then observation i is distributed to its left child; otherwise it is distributed to the right child. For every i , this distributing process iterates from the root node and ends in a leaf node. We use θ_k to denote the distribution parameters in the leaf nodes. For each node and a complete tree, their prior densities are:

$$\begin{aligned} [T_k] &= [s_k][c_k]^{s_k}[\mathbf{t}_k]^{s_k}[\theta_k]^{1-s_k} \\ [\mathbf{T}] &= [T_0] \prod_{k \in I_k} [T_{2k+1}][T_{2k+2}] \end{aligned} \quad (2)$$

For s_k, \mathbf{t}_k, c_k , we specify the prior distributions as follows:

$$\begin{aligned} s_k &\sim B(\exp(-\lfloor \log_2(k+1) \rfloor / \delta)) \\ c_k &\sim MN_m(\boldsymbol{\xi}), \text{ where } \boldsymbol{\xi} \sim Dir(\mathbf{1}_m) \\ [\mathbf{t}_k] &\propto 1 \end{aligned} \quad (3)$$

where B is Bernoulli distribution, MN_m is m -dimensional multinomial distribution. The hyper-parameter δ is the tuning parameter for which smaller δ results in smaller tree.

In each partition, objective priors are used for $\boldsymbol{\theta}$. If \mathbf{Y} is continuous, then $[\boldsymbol{\theta}_k] = [\mu, \sigma^2] \propto 1/\sigma^2$; if \mathbf{Y} is discrete, then $\boldsymbol{\theta}_k = \mathbf{p} \sim Dir(0.5 \cdot \mathbf{1})$. Note that Dir reduces to a *Beta* distribution when \mathbf{Y} is a binary outcome. To guarantee the posterior propriety of $\boldsymbol{\theta}$, we further require that each partition should have at least q observations and $q > 1$.

The posterior distribution of $\boldsymbol{\xi}$ reveals the proportion of instances that a certain variable is used in constructing a tree. One variable can be used more times than another, therefore resulting in a larger proportion in $\boldsymbol{\xi}$. Therefore, $\boldsymbol{\xi}$ can be utilized in variable selection and we name it as variable ranking probability.

3.2 Stick-Breaking Prior for \mathbf{W}

The changing dimension of the Dirichlet process creates difficulties in Bayesian sampling. Pioneering studies include exploring infinite state space with the reversible-jump Markov chain Monte Carlo (Green and Richardson, 2001) and with an auxiliary variable for possible new states (Neal, 2000). At the

same time, an equivalent construction named the stick-breaking process (Ishwaran and James, 2001) gained popularity for decreased computational burden. The stick-breaking process decomposes the Dirichlet process into an infinite series of *Beta* distributions:

$$\begin{aligned} w_1 &= v_1 \\ w_j &= v_j \prod_{k < j} (1 - v_k) \text{ for } j > 1 \end{aligned} \quad (4)$$

where each $v_j \stackrel{iid}{\sim} \text{Beta}(1, \alpha)$. This construction provides a straightforward illustration on the effects of adding/deleting a new cluster to/from the existing clusters.

Another difficulty in sampling is that j is infinite. Ishwaran and James (2001) proved that the $\max(j)$ can be truncated to 150 for a sample size of $n = 10^5$, and the results are indistinguishable from those obtained using larger numbers. Later, Kalli, Griffin, and Walker (2011) introduced the slice sampler, which avoids the approximate truncation. Briefly, the slice sampler adds a latent variable $u_i \sim U(0, 1)$ for each observation. The probability in (1) becomes:

$$[Y_i | \mathbf{X}_i, \{\mathbf{T}_j\}_j] = \sum_j^{\infty} \mathbf{1}(u_i < w_j) [Y_i | \mathbf{X}_i, \mathbf{T}_j] \quad (5)$$

due to $\int_0^1 \mathbf{1}(u_i < w_j) du_i = w_j$. The Monte Carlo sampling of u_i leads to omitting w_j 's that are too small. We found that the slice sampler usually leads to a smaller effective $\max j < 10$ for $n = 10^5$, hence more rapid convergence than a simple truncation.

4 Blocked Gibbs Sampling

We now explain the sampling algorithm for the BET model. Let $Z_i = j$ denote the latent assignment of the i th observation to the j th tree. Then the sampling scheme for the BET model involves iteration over two steps: tree growing and clustering.

4.1 Tree Growing: Updating $[\mathbf{T} | \mathbf{W}, \mathbf{Z}, \mathbf{Y}]$

Each tree with allocated data is grown in this step. We sample in the order of $[\mathbf{s}, \mathbf{c}, \mathbf{t}]$ and then $[\boldsymbol{\theta} | \mathbf{s}, \mathbf{c}, \mathbf{t}]$. As stated by Chipman, George, and McCulloch (1998), using $[\mathbf{Y} | \mathbf{s}, \mathbf{c}, \mathbf{t}]$ marginalized over $\boldsymbol{\theta}$ facilitates rapid change of the tree structure. After the tree is updated, the conditional sampling of $\boldsymbol{\theta}$

provides convenience in the next clustering step, where we compute $[Y_i|\mathbf{T}_j]$ for different j .

During the updating of $[\mathbf{s}, \mathbf{c}, \mathbf{t}]$, we found that using a random choice in grow/prune/swap/change (GPSC) in one Metropolis-Hasting (MH) step (Chipman, George, and McCulloch, 1998), is not sufficient to grow large trees for our model. This is not a drawback of the proposal mechanism, but is instead primarily due to the notion that following this clustering process would distribute the data entries to many small trees, if any large tree has not yet formed. In other words, the goal is to have our model prioritize “first in growing the tree, second in clustering” instead of the other order.

Therefore, we devise a new Gibbs sampling scheme, which sequentially samples the full conditional distribution $[s_k|(s_k)]$, $[c_k|(c_k)]$ and $[t_k|(t_k)]$. For each update, MH criterion is used. We restrict updates of \mathbf{c}, \mathbf{t} that result in an empty node such that \mathbf{s} do not change in these steps. The major difference in this approach compared to the GPSC method is that, rather than one random change in one random node, we use micro steps to exhaustively explore possible changes in every node, which increases chain convergence.

Besides increasing the convergence rate, the other function of the Gibbs sampler is to force each change in the tree structure to be small and local. Although some radical change steps (Wu, Tjelmeland, and West, 2007; Pratola, 2013) can facilitate the jumps between the modes in a single tree, for a mixture of trees, local changes and mode sticking are useful to prevent label switching.

4.2 Clustering: Updating $[\mathbf{W}, \mathbf{Z}|\mathbf{T}, \mathbf{Y}]$

In this step, we take advantage of the latent uniform variable \mathbf{U} in the slice sampler. In order to sample from the joint density $[\mathbf{U}, \mathbf{W}, \mathbf{Z}|\mathbf{T}]$, we use the blocked Gibbs sampling again, in the order of the following marginal densities $[\mathbf{W}|\mathbf{Z}, \mathbf{T}]$, $[\mathbf{U}|\mathbf{W}, \mathbf{Z}, \mathbf{T}]$ and $[\mathbf{Z}|\mathbf{W}, \mathbf{U}, \mathbf{T}]$.

$$\begin{aligned} v_j &\sim \text{Beta}(1 + \sum_i \mathbf{1}(Z_i = j), \alpha + \sum_i \mathbf{1}(Z_i > j)) \\ w_j &= v_j \prod_{k < j} (1 - v_k) \\ u_i &\sim U(0, w_{Z_i}) \end{aligned} \tag{6}$$

During the following update of the assignment \mathbf{Z} , two variables simplify the computation: the truncation effects of the u_i ’s keep the number of candidate states finite; the values of $\boldsymbol{\theta}$ in each tree \mathbf{T} circumvent the need to recompute the marginal likelihood. The probability of assigning the i th observation to the j th tree is

$$[Z_i = j|\mathbf{W}, \mathbf{Y}, \mathbf{T}_j] \propto \mathbf{1}(w_j > u_i)[Y_i|\mathbf{T}_j] \tag{7}$$

4.3 Posterior Inference: Choosing the Best Ensemble of Trees

In the posterior analyses of the Markov chain, we could consider the marginal likelihood $\prod_i \int_{Z_i} [y_i | T_{Z_i}] dP(Z_i)$, but it would involve costly allocation of the data over all the candidate trees. Therefore, we use the joint likelihood with the tree assignment $\prod_i [y_i | T_{Z_i}] [Z_i]$ as the criterion for choosing the best ensemble of trees.

For prediction purposes, we define two types of estimators: a cluster-specific estimator and an ensemble estimator. The former is defined as $\mathbb{E}(\theta | T_{Z_i})$, where assignment Z_i is known. The latter is defined as $\sum_j w_j \mathbb{E}(\theta | T_{Z_j})$, where assignment Z_i is unknown. In theory, conditional on the collected posterior samples being the same, the cluster-specific estimator has smaller variance than the ensemble estimator. In practice, the latter is more applicable since we often do not know Z_i for a new observation until Y_i is observed.

5 Simulation Studies

In this section, we demonstrate the clustering capability of BET through simulations of three scenarios: (1) single way of partitioning in \mathbf{X} with uni-modal $[\mathbf{Y} | \mathbf{X}]$; (2) single way of partitioning in \mathbf{X} with multi-modal $[\mathbf{Y} | \mathbf{X}]$; (3) multiple ways of partitioning of \mathbf{X} with possible multi-modal $[\mathbf{Y} | \mathbf{X}]$. For case (1), the model should converge to one cluster, whereas cases (2) and (3) should both result in multiple clusters.

5.1 Unique Partitioning Scheme with Unimodal Likelihood

This is the simplest yet still likely scenario, in which the data is homogeneous and can be easily classified with a unique partition. We adopt the simulation setting used by Wu, Tjelmeland, and West (2007) and generate 300 test samples as shown in Table 1. In this dataset, three partitions can be classified based on combinations of either $\{X_1, X_2\}$ or $\{X_2, X_3\}$, where two choices are equivalent.

We ran the BET model on the data for 10,000 steps and used the last 5,000 steps as the results. The best tree is shown in Figure 1. Three partitions are correctly identified. The choice of X and splitting values are consistent with the simulation parameters. More importantly, the majority of observations are clustered together in a single tree. In each partition, the number of observations and the parameter estimates are very close to the true values. This simulation shows that BET functions as a simple Bayesian CART method when there is no need of clustering.

Index	X_1	X_2	X_3	Y
1...100	U(0.1,0.4)	U(0.1,0.4)	U(0.6,0.9)	N(1.0,0.5 ²)
101...200	U(0.1,0.4)	U(0.6,0.9)	U(0.6,0.9)	N(3.0,0.5 ²)
201...300	U(0.6,0.9)	U(0.1,0.9)	U(0.1,0.4)	N(5.0,0.5 ²)

Table 1: Setting for Simulation Study I

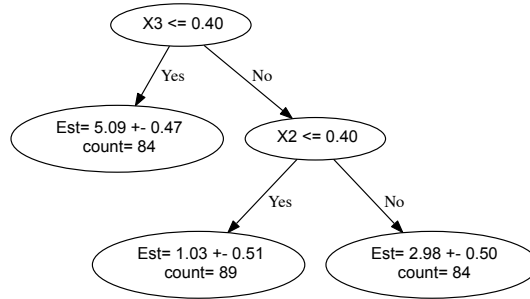


Figure 1: Simulation Study I shows one cluster is found and the partitioning scheme is correctly uncovered

5.2 Unique Partitioning Scheme with Multi-modal Likelihood

This example is to illustrate the scenario of having a “mixture distribution inside one tree”. We duplicate the previous observations in each partition and change one set of their means. As shown in Table 2, each partition now becomes bimodal. Since partitioning based on $\{X_1, X_3\}$ is equivalent to $\{X_1, X_2\}$, we drop X_3 for clearer demonstration.

The results are shown in Figure 2. Two clusters are correctly identified by BET. The fitted splitting criteria and the estimated parameters are consistent with the true values. It is interesting to note that in the leftmost nodes of the two clusters, there is a significant overlap in distribution (within one standard deviation of normal means). As a result, compared with the original counts in the data generation, some randomness is observed in these two nodes. Nevertheless, the two fitted trees are almost the same as anticipated.

Index	X_1	X_2	Y
1...100	U(0.1,0.4)	U(0.1,0.4)	N(1.0,0.5 ²)
101...200	U(0.1,0.4)	U(0.6,0.9)	N(3.0,0.5 ²)
201...300	U(0.6,0.9)	U(0.1,0.9)	N(5.0,0.5 ²)
301...400	U(0.1,0.4)	U(0.1,0.4)	N(1.5,0.5 ²)
401...500	U(0.1,0.4)	U(0.6,0.9)	N(5.5,0.5 ²)
501...600	U(0.6,0.9)	U(0.1,0.9)	N(3.5,0.5 ²)

Table 2: Setting for Simulation Study II

5.3 Mixed Partitioning Scheme with Unimodal or Multimodal Likelihood

This scenario reflects the most complicated case, which is quite realistic in large, heterogeneous data. We again duplicate the data in Table 1 and then change the partition scheme for indices 301...600. This further increases the challenge as shown in Figure 3. The data are now formed by a mixed partitioning scheme, which has both overlapping and unique regions. Their distribution means are labeled in each region. It is worth noting that the shared partition in the upper left is bimodal.

Index	X_1	X_2	Y
1...100	U(0.1,0.4)	U(0.1,0.4)	N(1.0,0.5 ²)
101...200	U(0.1,0.4)	U(0.6,0.9)	N(3.0,0.5 ²)
201...300	U(0.6,0.9)	U(0.1,0.9)	N(5.0,0.5 ²)
301...400	U(0.1,0.4)	U(0.6,0.9)	N(5.0,0.5 ²)
401...500	U(0.6,0.9)	U(0.6,0.9)	N(1.0,0.5 ²)
501...600	U(0.1,0.9)	U(0.1,0.4)	N(3.0,0.5 ²)

Table 3: Setting for Simulation Study III

To ensure convergence, we ran the model for 20,000 steps and discard the first 10,000 as burn-in steps. Among the total 10,000 steps, the step numbers that correspond to 1, 2 and 3 clusters are 214, 9665 and 121. Clearly, the 2-cluster is the most probable model for the data.

The best ensemble of trees are shown in Figure 4. This is consistent to the data generation diagram in Figure 3, since the two means in the upper left region are exchangeable due to its bimodal nature.

These simulation results show that BET is capable of detecting not only

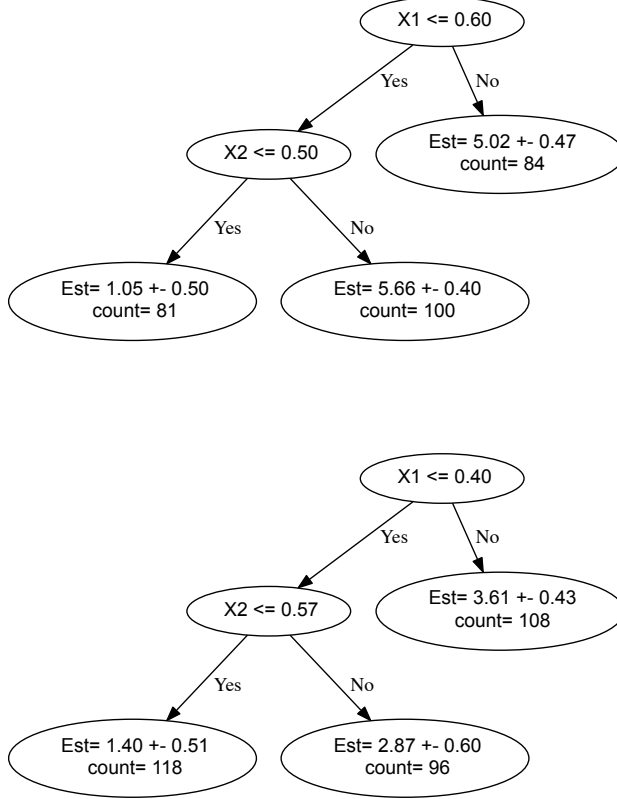


Figure 2: Simulation Study II shows two clusters are found and parameters are correctly estimated

the mixing of distributions, but also the mixing of trees. In the following sections, we test its prediction performance through real data examples.

6 Breast Cancer Data Example

We first demonstrate the performance of the Gibbs sampler in the BET model with homogenous data. We examined breast cancer data available from the machine learning repository of the University of California at Irvine (Bache and Lichman, 2013). This data was originally the work of Wolberg and Mangasarian (1990) and has been subsequently utilized in numerous

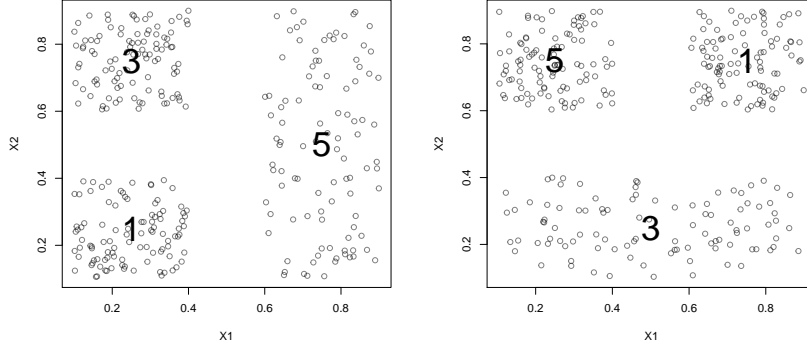


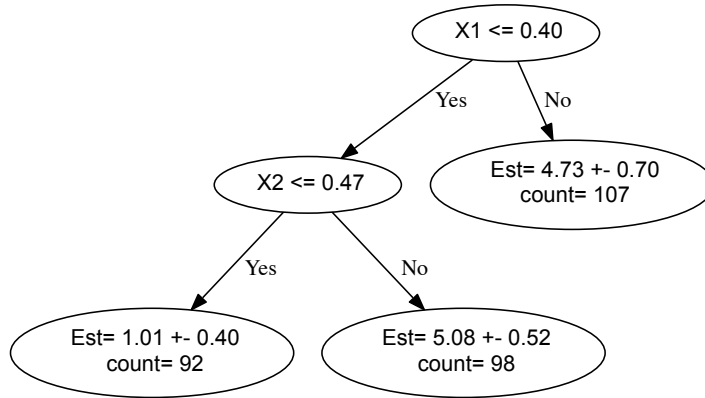
Figure 3: Simulation Study III has data from different partitioning schemes mixed together. The means are labeled in the center of each region. The shared region in the upper left has mixed means.

application studies. We focus on the outcome of breast cancer as the response variable. We define benign and malignant results as $y = 0$ and $y = 1$, respectively. Our application includes the nine clinical covariates, which are clump thickness, uniformity of cell size, uniformity of cell shape, marginal adhesion, single epithelial cell size, bare nuclei, bland chromatin, normal nucleoli and mitoses. We consider data for 683 females who have no missing values for the outcome and covariates at hand. We ran the model for 110,000 steps and discarded the first 10,000 steps.

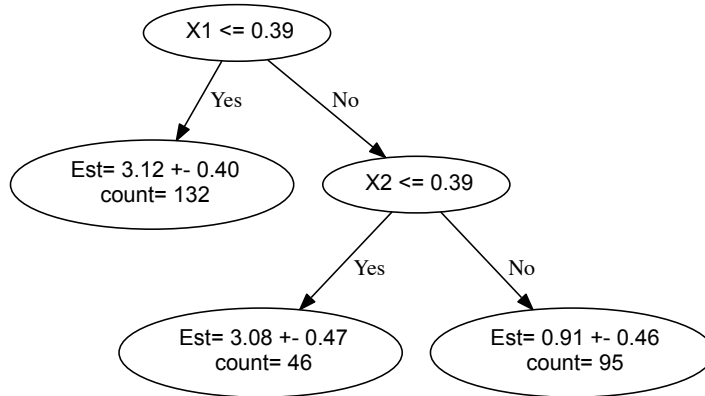
The results show the chain converges to the joint log-likelihood $[Y, Z|T]$ at mean -138.3 and the conditional log-likelihood $[Y|Z, T]$ at mean -53.8 (Figure 5(a)). The second number is slightly improved over the result reported by Wu, Tjelmeland, and West (2007).

We use 0.5 as the cut-off point of estimates to dichotomize 0 and 1, then use the proportion of incorrect classification as the misclassification rate (MCR). The majority of the collected steps correspond to only one cluster (Figure 5(c)). This suggests the data is highly homogeneous, which contributes to the low misclassification rate (Figure 5(b)). The MCR is 0.025 ± 0.007 (Figure 5(c)), with the minimum equal to 0.013.

To assess prediction performance, we carried out half-split cross-validation. We randomly selected 342 observations as the training set. Predictions are made on the remaining 341 observations. We use the ensemble estimator for the MCR calculation. We repeated the test for ten separate times, each with a different seed for the random split. The resulting MCR is as low as 0.036 ± 0.006 , which is close to the one reported by Wu, Tjelmeland, and West (2007). These results support that the Gibbs sampling algorithm can rapidly move the chain in the tree structure space \mathcal{T} .




(a) Cluster 1




(b) Cluster 2

Figure 4: Simulation Study III shows two clusters correctly identified

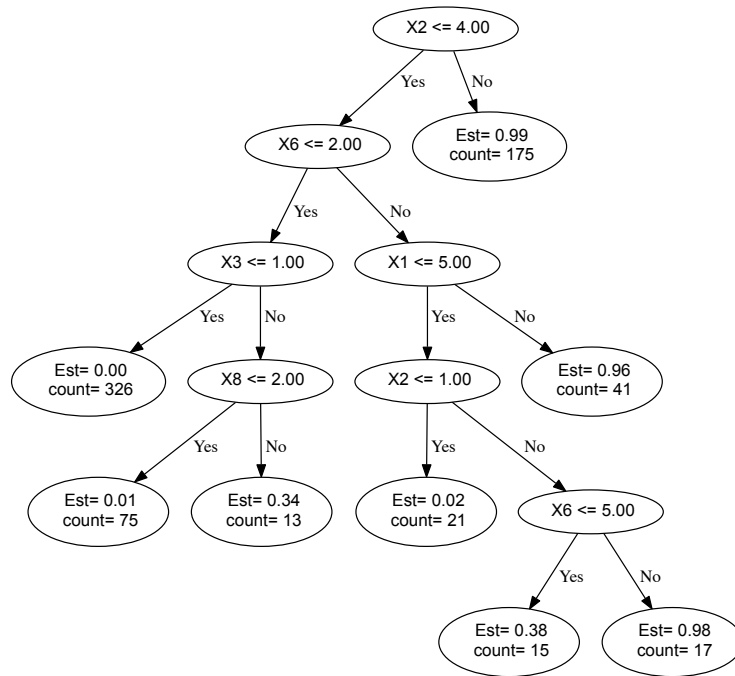
Next we test BET on a large dataset, which possibly contains heterogeneous data and simultaneously illustrate its regression performance.



(a) Log-likelihood given assignment



(b) Misclassification rate with cluster-specific estimator



(c) Best ensemble: only one cluster is found

Figure 5: Result of breast cancer test

7 Cystic Fibrosis Data Example

We used lung function data obtained from the Cystic Fibrosis Foundation Patent Registry (Cystic Fibrosis Foundation, 2012). Percent predicted of forced expiratory volume in 1 second ($FEV_1\%$) is a continuous measure of lung function in cystic fibrosis (CF) patients obtained at each clinical visit. We have previously demonstrated that the rates of $FEV_1\%$ change non-linearly (Szczesniak, McPhail, Duan, Macaluso, Amin, and Clancy, 2013), which can be described via semiparametric regression using penalized cubic splines. Although trees may not outperform spline methods in prediction of continuous outcomes, they provide reliable information for variable selection when traditional methods such as p-value inspection and AIC may fall short.

We used longitudinal data from 3,500 subjects (a total of 60,000 entries) from the CF dataset and utilized eight clinical covariates: baseline $FEV_1\%$, age, gender, infections (each abbreviated as MRSA, PA, BC, CFRD) and insurance (a measure of socioeconomic status, SES). We randomly selected longitudinal data from 2,943 subjects (roughly 50,000 entries) and used these data as the training set. We then carried out prediction on the remaining 10,000 entries.

We illustrate the prediction results of BET in Figure 6(a and b). With the assumption of one constant mean per partition, the predicted line takes the shape of step functions, which correctly captured the declining trend of $FEV_1\%$. The prediction seems unbiased, as the difference between the predicted and true values are symmetric about the diagonal line. We also computed the difference metrics with the true values. The results are listed in Table 4.

Besides the BET model, we also tested other popular regression tree methods (with corresponding R packages), such as CART (“rpart”, Therneau, Atkinson, et al. (1997)), Random Forests (“randomForest” Breiman (2001)), Boosting (“gbm” Friedman (2001)) and BART (“bartMachine” Chipman, George, McCulloch, et al. (2010)). Since the tested data are essentially longitudinal data, we can choose whether to group observations by subjects or by entries alone. In the subject-clustered version, we first used one outcome entry of each subject in the prediction subset, computed the most likely cluster and then computed the cluster-specific predictor; in the entry-clustered version, we simply used the ensemble predictor. We did not see an obvious difference between the two predictors. This is likely due to entry-clustered BET achieving better fit, which compensates for less accuracy of the ensemble predictor.

In the comparison among the listed tree-based methods, the two Bayesian methods BET and BART provide the closest results to spline regression in prediction accuracy. Similar to the relation between BART and Boosting, BET can be viewed as the Bayes counterpart of Random Forests. Besides the use of prior, one important distinction is that the Random Forests approach uses the average of bootstrap sample estimates, whereas BET uses


Model	RMSE	MAD
Spline Regression	16.60	10.07
CART	18.07	10.32
Random Forests (5 trees)	17.38	11.29
Random Forests (50 trees)	17.13	11.28
Boosting (1000 trees)	20.34	14.22
BART (50 trees)	16.72	10.32
BET (clustered by subjects)	16.97	10.57
BET (clustered by entries)	16.70	10.13

Table 4: Cross-validation results with various methods applied on cystic fibrosis data.


the weighted average of cluster sample estimates. In Random Forests, the number of bootstrap sample needs to be specified by the user; while in BET, it is determined by the data through Dirichlet process. During this test, Random Forests used 50 trees; while BET converged to only 2-3 trees (Figure 6(c)) and achieved similar prediction accuracy. The tree structures are shown in the appendix.

Lastly, we focus on the variable selection issue. Although some information criteria have been established in this area, such as AIC, BIC and Bayes Factor, their mechanisms are complicated by the necessity to fit different models to the data for multiple times. The inclusion probability (reviewed by O’Hara, Sillanpää, et al. (2009)) is an attractive alternative, which penalizes the addition of more variables through the inclusion prior. In the tree-based method, however, since multiplicity is not a concern, it is possible to compare several variables of similar importances at the same time, without inclusion or exclusion.


Since multiple trees are present, we use the weighted posterior $\bar{\xi} = \sum_j w_j \xi_j$ as the measure for variable importance. We plot the variable ranking probability $\bar{\xi}$ for each covariate (Figure 6). The interpretation of this value follows naturally as how likely a covariate is chosen in forming the trees. Therefore, the ranking of $\bar{\xi}$ reveals the order of importance of the covariates. This concept is quite similar to the variable importance measure invented in Random Forests. The difference is that their index is ranked in the decrease of accuracy after permuting of a covariate; ours is purely a probabilistic measure. Regardless of this difference, the ranking of variables from the two models are remarkably similar (see Supplementary Information for results of Random Forests): baseline FEV₁% and age are the two most important variables while gender and MRSA seem to play the least important roles.




(a) True values (dots) and the 95% credible interval of the BET predicted means: the estimates are piecewise constant in a single subject.



(b) True values and the BET predicted means: the estimates are piecewise constant and seem unbiased.



(c) Number of clusters found



(d) The variable selection probabilities

Figure 6: Results of FEV₁% fitting and prediction.

8 Discussion

Empirically, compared with a single model, a group of models (an ensemble) usually has better performance in inference and prediction. This view is also supported in the field of the decision trees. As reviewed by Hastie, Tibshirani, Friedman, Hastie, Friedman, and Tibshirani (2009), the use of decision trees has benefited from multiple-tree methods such as Random Forests, Boosting and Bayesian tree averaging. One interesting question that remains, however, is how many models are enough, or, how many are more than enough?

To address this issue, machine learning algorithms usually resort to repeated tests of cross validation or out-of-bootstrap error calculation, in which ensemble size is gradually increased until performance starts to degrade. In Bayesian model averaging, the number of steps to keep is often “the more, the better”, as long as the steps have low autocorrelation.

Our proposed method, BET, demonstrates a more efficient way to create an ensemble with the same predictive capability but much smaller size. With the help of the Dirichlet process, the self-reinforcing behavior of large clusters reduces the number of needed clusters (sub-models). Rather than using a simple average over many trees, we showed that using a weighted average over a few important trees can provide the same accuracy.

It is worth comparing BET with other mixture models. In the latter, the component distributions are typically continuous and unimodal; in BET, each component tree is discrete, and more importantly, multi-modal itself. This construction could have created caveats in model fitting, as one can imagine only obtaining a large ensemble of very small trees. We circumvented this issue by applying Gibbs sampling in each tree, which rapidly increases the fit of tree to the data during tree growing, and decreases the chance that they are scattered to more clusters.

It is also of interest to develop an empirical algorithm for BET. One possible extension is to use a local optimization technique (also known as “greedy algorithm”) under some randomness to explore the tree structure. This implementation may be not difficult, since users can access existing CART packages to grow trees for subsets of data, then update clustering as mentioned previously.

Acknowledgments

Partial funding was provided by the Cystic Fibrosis Foundation Research and Development Program (grant number R457-CR11). The authors are grateful to the Cystic Fibrosis Foundation Patient Registry Committee for their thoughtful comments and data dispensation. The breast cancer domain was obtained from the University Medical Centre, Institute of Oncology, Ljubljana, Yugoslavia. The authors thank M. Zwitter and M. Soklic for availability of this data.

References

- K. Bache and M. Lichman. UCI machine learning repository, 2013. URL <http://archive.ics.uci.edu/ml>.
- L. Breiman, J. Friedman, R. Olshen, and C. Stone. *Classification and Regression Trees*. Wadsworth and Brooks, Monterey, CA, 1984.
- Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.
- Hugh A Chipman, Edward I George, and Robert E McCulloch. Bayesian cart model search. *Journal of the American Statistical Association*, 93(443):935–948, 1998.
- Hugh A Chipman, Edward I George, Robert E McCulloch, et al. Bart: Bayesian additive regression trees. *The Annals of Applied Statistics*, 4(1): 266–298, 2010.
- 1 Cystic Fibrosis Foundation. Cystic fibrosis foundation patient registry 2012 annual data report. 2012.
- David GT Denison, Bani K Mallick, and Adrian FM Smith. A bayesian cart algorithm. *Biometrika*, 85(2):363–377, 1998.
- Jerome H Friedman. Greedy function approximation: a gradient boosting machine. *Annals of Statistics*, pages 1189–1232, 2001.
- Jerome H Friedman. Stochastic gradient boosting. *Computational Statistics & Data Analysis*, 38(4):367–378, 2002.
- Peter J Green and Sylvia Richardson. Modelling heterogeneity with and without the dirichlet process. *Scandinavian journal of statistics*, 28(2): 355–375, 2001.
- Trevor Hastie, Robert Tibshirani, Jerome Friedman, T Hastie, J Friedman, and R Tibshirani. *The elements of statistical learning*, volume 2. Springer, 2009.
- Hemant Ishwaran and Lancelot F. James. Gibbs sampling methods for stick-breaking priors. *Journal of the American Statistical Association*, 96(453): 161–173, 2001.
- Maria Kalli, Jim E Griffin, and Stephen G Walker. Slice sampling mixture models. *Statistics and computing*, 21(1):93–105, 2011.
- Radford M Neal. Markov chain sampling methods for dirichlet process mixture models. *Journal of computational and graphical statistics*, 9(2):249–265, 2000.

- Robert B O'Hara, Mikko J Sillanpää, et al. A review of bayesian variable selection methods: what, how and which. *Bayesian analysis*, 4(1):85–117, 2009.
- MT Pratola. Efficient metropolis-hastings proposal mechanisms for bayesian regression tree models. *arXiv preprint arXiv:1312.1895*, 2013.
- Rhonda D Szczesniak, Gary L McPhail, Leo L Duan, Maurizio Macaluso, Raouf S Amin, and John P Clancy. A semiparametric approach to estimate rapid lung function decline in cystic fibrosis. *Annals of epidemiology*, 23(12):771–777, 2013.
- Terry M Therneau, Elizabeth J Atkinson, et al. An introduction to recursive partitioning using the rpart routines. 1997.
- William H Wolberg and Olvi L Mangasarian. Multisurface method of pattern separation for medical diagnosis applied to breast cytology. *Proceedings of the national academy of sciences*, 87(23):9193–9196, 1990.
- Yuhong Wu, Håkon Tjelmeland, and Mike West. Bayesian cart: Prior specification and posterior simulation. *Journal of Computational and Graphical Statistics*, 16(1):44–66, 2007.

Supplementary Information

Variable Importance calculated by Random Forests with 50 trees, using cystic fibrosis data.




Figure 7: Variable importance calculated by Random Forests with 50 trees




Figure 8: Cluster 1 found by the BET model applied on the cystic fibrosis data




Figure 9: Cluster 2 found by the BET model applied on the cystic fibrosis data